

SPECIFICATION
(Sprint Docket No. 1732)

TO ALL WHOM IT MAY CONCERN:

5 Be it known that we, **Kevin PACKINGHAM**, a citizen of the United States and resident of Overland Park, Kansas, and **Robert W. HAMMOND**, a citizen of the United States and resident of Overland Park, Kansas, have invented a new and useful:

METHOD AND SYSTEM FOR
CONSOLIDATED MESSAGE NOTIFICATION
IN A VOICE COMMAND PLATFORM

10 the following of which is a specification.

09964099-092604

RELATED APPLICATIONS

This application is related to the following commonly owned applications filed on the same date: (i) "Method and System for Dynamic Control Over Voice-Processing in a Voice Command Platform," naming Kevin Packingham, Elizabeth Roche and Terry T. Yu as co-inventors, (ii) "Method and System for Enhanced Response to Voice Commands in a Voice Command Platform," naming Kevin Packingham and Elizabeth Roche as co-inventors, and (iii) "Method and System for Use of Navigation History in a Voice Command Platform," naming Kevin Packingham as inventor. The entirety of each of these other applications is hereby incorporated by reference.

BACKGROUND

1. Field of the Invention

The present invention relates to telecommunications and, more particularly, to functions of a voice command platform.

2. Description of Related Art

A voice command platform provides an interface between speech communication with a user and computer-executed voice command applications. Generally speaking, a person can call a voice command platform from any telephone and, by speaking commands, can browse through navigation points (e.g., applications and/or menu items within the applications) to access and communicate information. The voice command platform can thus receive spoken commands from the user and use the commands to guide its execution of voice command applications, and the voice command platform can "speak" to a user as dictated by logic in voice command applications.

For instance, a person may call a voice command platform, and the platform may apply a voice command application that causes the platform to speak to the user, "Hello. Would you like to hear a weather forecast, sports scores, or stock quotes?" In response, the person may state to the platform, "weather forecast." Given this response, the application may cause the platform to load and execute a subsidiary weather forecasting application. The weather forecasting application may direct the platform to speak another speech prompt to the person, such as "Would you like to hear today's weather or an extended forecast?" The person may then respond, and the weather forecasting application may direct the voice command platform to execute additional logic or to load and execute another application based on the person's response.

A robust voice command platform should therefore be able to (i) receive and recognize speech spoken by a user and (ii) provide speech to a user. The platform can achieve these functions in various ways.

On the incoming side, for instance, the platform may include an analog-to-digital (A-D) converter for converting an analog speech signal from a user into a digitized incoming speech signal. (Alternatively, the user's speech signal might already be digitized, as in a voice-over-IP communication system, for instance, in which case A-D conversion would be unnecessary). The platform may then include a speech recognition (SR) engine, which functions to analyze the digitized incoming speech signal and to identify words in the speech. The SR engine will typically be a software module executable by a processor.

Usually, a voice command application will specify which words or "grammars" a user can speak in response to a prompt for instance. Therefore, the SR engine will seek to identify

one of the possible spoken responses. (Alternatively, the SR engine may operate to identify any words without limitation).

5 In order to identify words in the incoming speech, the SR engine will typically include or have access to a dictionary database of "phonemes" (small units of speech that distinguish one utterance from another). The SR engine will then analyze the waveform represented by the incoming digitized speech signal and, based on the dictionary database, will determine whether the waveform represents particular words. For instance, if a voice command application allows for a user to respond to a prompt with the grammars "sales," "service" or "operator", the SR engine may identify the sequence of one or more phonemes that makes up each of these grammars respectively. The SR engine may then analyze the waveform of the incoming digitized speech signal in search of a waveform that represents one of those sequences of phonemes. Once the SR engine finds a match, the voice command platform may continue processing the application in view of the user's spoken response.

10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
105
110
115
120
125
130
135
140
145
150
155
160
165
170
175
180
185
190
195
200
205
210
215
220
225
230
235
240
245
250
255
260
265
270
275
280
285
290
295
300
305
310
315
320
325
330
335
340
345
350
355
360
365
370
375
380
385
390
395
400
405
410
415
420
425
430
435
440
445
450
455
460
465
470
475
480
485
490
495
500
505
510
515
520
525
530
535
540
545
550
555
560
565
570
575
580
585
590
595
600
605
610
615
620
625
630
635
640
645
650
655
660
665
670
675
680
685
690
695
700
705
710
715
720
725
730
735
740
745
750
755
760
765
770
775
780
785
790
795
800
805
810
815
820
825
830
835
840
845
850
855
860
865
870
875
880
885
890
895
900
905
910
915
920
925
930
935
940
945
950
955
960
965
970
975
980
985
990
995
1000
1005
1010
1015
1020
1025
1030
1035
1040
1045
1050
1055
1060
1065
1070
1075
1080
1085
1090
1095
1100
1105
1110
1115
1120
1125
1130
1135
1140
1145
1150
1155
1160
1165
1170
1175
1180
1185
1190
1195
1200
1205
1210
1215
1220
1225
1230
1235
1240
1245
1250
1255
1260
1265
1270
1275
1280
1285
1290
1295
1300
1305
1310
1315
1320
1325
1330
1335
1340
1345
1350
1355
1360
1365
1370
1375
1380
1385
1390
1395
1400
1405
1410
1415
1420
1425
1430
1435
1440
1445
1450
1455
1460
1465
1470
1475
1480
1485
1490
1495
1500
1505
1510
1515
1520
1525
1530
1535
1540
1545
1550
1555
1560
1565
1570
1575
1580
1585
1590
1595
1600
1605
1610
1615
1620
1625
1630
1635
1640
1645
1650
1655
1660
1665
1670
1675
1680
1685
1690
1695
1700
1705
1710
1715
1720
1725
1730
1735
1740
1745
1750
1755
1760
1765
1770
1775
1780
1785
1790
1795
1800
1805
1810
1815
1820
1825
1830
1835
1840
1845
1850
1855
1860
1865
1870
1875
1880
1885
1890
1895
1900
1905
1910
1915
1920
1925
1930
1935
1940
1945
1950
1955
1960
1965
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
2025
2030
2035
2040
2045
2050
2055
2060
2065
2070
2075
2080
2085
2090
2095
2100
2105
2110
2115
2120
2125
2130
2135
2140
2145
2150
2155
2160
2165
2170
2175
2180
2185
2190
2195
2200
2205
2210
2215
2220
2225
2230
2235
2240
2245
2250
2255
2260
2265
2270
2275
2280
2285
2290
2295
2300
2305
2310
2315
2320
2325
2330
2335
2340
2345
2350
2355
2360
2365
2370
2375
2380
2385
2390
2395
2400
2405
2410
2415
2420
2425
2430
2435
2440
2445
2450
2455
2460
2465
2470
2475
2480
2485
2490
2495
2500
2505
2510
2515
2520
2525
2530
2535
2540
2545
2550
2555
2560
2565
2570
2575
2580
2585
2590
2595
2600
2605
2610
2615
2620
2625
2630
2635
2640
2645
2650
2655
2660
2665
2670
2675
2680
2685
2690
2695
2700
2705
2710
2715
2720
2725
2730
2735
2740
2745
2750
2755
2760
2765
2770
2775
2780
2785
2790
2795
2800
2805
2810
2815
2820
2825
2830
2835
2840
2845
2850
2855
2860
2865
2870
2875
2880
2885
2890
2895
2900
2905
2910
2915
2920
2925
2930
2935
2940
2945
2950
2955
2960
2965
2970
2975
2980
2985
2990
2995
3000
3005
3010
3015
3020
3025
3030
3035
3040
3045
3050
3055
3060
3065
3070
3075
3080
3085
3090
3095
3100
3105
3110
3115
3120
3125
3130
3135
3140
3145
3150
3155
3160
3165
3170
3175
3180
3185
3190
3195
3200
3205
3210
3215
3220
3225
3230
3235
3240
3245
3250
3255
3260
3265
3270
3275
3280
3285
3290
3295
3300
3305
3310
3315
3320
3325
3330
3335
3340
3345
3350
3355
3360
3365
3370
3375
3380
3385
3390
3395
3400
3405
3410
3415
3420
3425
3430
3435
3440
3445
3450
3455
3460
3465
3470
3475
3480
3485
3490
3495
3500
3505
3510
3515
3520
3525
3530
3535
3540
3545
3550
3555
3560
3565
3570
3575
3580
3585
3590
3595
3600
3605
3610
3615
3620
3625
3630
3635
3640
3645
3650
3655
3660
3665
3670
3675
3680
3685
3690
3695
3700
3705
3710
3715
3720
3725
3730
3735
3740
3745
3750
3755
3760
3765
3770
3775
3780
3785
3790
3795
3800
3805
3810
3815
3820
3825
3830
3835
3840
3845
3850
3855
3860
3865
3870
3875
3880
3885
3890
3895
3900
3905
3910
3915
3920
3925
3930
3935
3940
3945
3950
3955
3960
3965
3970
3975
3980
3985
3990
3995
4000
4005
4010
4015
4020
4025
4030
4035
4040
4045
4050
4055
4060
4065
4070
4075
4080
4085
4090
4095
4100
4105
4110
4115
4120
4125
4130
4135
4140
4145
4150
4155
4160
4165
4170
4175
4180
4185
4190
4195
4200
4205
4210
4215
4220
4225
4230
4235
4240
4245
4250
4255
4260
4265
4270
4275
4280
4285
4290
4295
4300
4305
4310
4315
4320
4325
4330
4335
4340
4345
4350
4355
4360
4365
4370
4375
4380
4385
4390
4395
4400
4405
4410
4415
4420
4425
4430
4435
4440
4445
4450
4455
4460
4465
4470
4475
4480
4485
4490
4495
4500
4505
4510
4515
4520
4525
4530
4535
4540
4545
4550
4555
4560
4565
4570
4575
4580
4585
4590
4595
4600
4605
4610
4615
4620
4625
4630
4635
4640
4645
4650
4655
4660
4665
4670
4675
4680
4685
4690
4695
4700
4705
4710
4715
4720
4725
4730
4735
4740
4745
4750
4755
4760
4765
4770
4775
4780
4785
4790
4795
4800
4805
4810
4815
4820
4825
4830
4835
4840
4845
4850
4855
4860
4865
4870
4875
4880
4885
4890
4895
4900
4905
4910
4915
4920
4925
4930
4935
4940
4945
4950
4955
4960
4965
4970
4975
4980
4985
4990
4995
5000
5005
5010
5015
5020
5025
5030
5035
5040
5045
5050
5055
5060
5065
5070
5075
5080
5085
5090
5095
5100
5105
5110
5115
5120
5125
5130
5135
5140
5145
5150
5155
5160
5165
5170
5175
5180
5185
5190
5195
5200
5205
5210
5215
5220
5225
5230
5235
5240
5245
5250
5255
5260
5265
5270
5275
5280
5285
5290
5295
5300
5305
5310
5315
5320
5325
5330
5335
5340
5345
5350
5355
5360
5365
5370
5375
5380
5385
5390
5395
5400
5405
5410
5415
5420
5425
5430
5435
5440
5445
5450
5455
5460
5465
5470
5475
5480
5485
5490
5495
5500
5505
5510
5515
5520
5525
5530
5535
5540
5545
5550
5555
5560
5565
5570
5575
5580
5585
5590
5595
5600
5605
5610
5615
5620
5625
5630
5635
5640
5645
5650
5655
5660
5665
5670
5675
5680
5685
5690
5695
5700
5705
5710
5715
5720
5725
5730
5735
5740
5745
5750
5755
5760
5765
5770
5775
5780
5785
5790
5795
5800
5805
5810
5815
5820
5825
5830
5835
5840
5845
5850
5855
5860
5865
5870
5875
5880
5885
5890
5895
5900
5905
5910
5915
5920
5925
5930
5935
5940
5945
5950
5955
5960
5965
5970
5975
5980
5985
5990
5995
6000
6005
6010
6015
6020
6025
6030
6035
6040
6045
6050
6055
6060
6065
6070
6075
6080
6085
6090
6095
6100
6105
6110
6115
6120
6125
6130
6135
6140
6145
6150
6155
6160
6165
6170
6175
6180
6185
6190
6195
6200
6205
6210
6215
6220
6225
6230
6235
6240
6245
6250
6255
6260
6265
6270
6275
6280
6285
6290
6295
6300
6305
6310
6315
6320
6325
6330
6335
6340
6345
6350
6355
6360
6365
6370
6375
6380
6385
6390
6395
6400
6405
6410
6415
6420
6425
6430
6435
6440
6445
6450
6455
6460
6465
6470
6475
6480
6485
6490
6495
6500
6505
6510
6515
6520
6525
6530
6535
6540
6545
6550
6555
6560
6565
6570
6575
6580
6585
6590
6595
6600
6605
6610
6615
6620
6625
6630
6635
6640
6645
6650
6655
6660
6665
6670
6675
6680
6685
6690
6695
6700
6705
6710
6715
6720
6725
6730
6735
6740
6745
6750
6755
6760
6765
6770
6775
6780
6785
6790
6795
6800
6805
6810
6815
6820
6825
6830
6835
6840
6845
6850
6855
6860
6865
6870
6875
6880
6885
6890
6895
6900
6905
6910
6915
6920
6925
6930
6935
6940
6945
6950
6955
6960
6965
6970
6975
6980
6985
6990
6995
7000
7005
7010
7015
7020
7025
7030
7035
7040
7045
7050
7055
7060
7065
7070
7075
7080
7085
7090
7095
7100
7105
7110
7115
7120
7125
7130
7135
7140
7145
7150
7155
7160
7165
7170
7175
7180
7185
7190
7195
7200
7205
7210
7215
7220
7225
7230
7235
7240
7245
7250
7255
7260
7265
7270
7275
7280
7285
7290
7295
7300
7305
7310
7315
7320
7325
7330
7335
7340
7345
7350
7355
7360
7365
7370
7375
7380
7385
7390
7395
7400
7405
7410
7415
7420
7425
7430
7435
7440
7445
7450
7455
7460
7465
7470
7475
7480
7485
7490
7495
7500
7505
7510
7515
7520
7525
7530
7535
7540
7545
7550
7555
7560
7565
7570
7575
7580
7585
7590
7595
7600
7605
7610
7615
7620
7625
7630
7635
7640
7645
7650
7655
7660
7665
7670
7675
7680
7685
7690
7695
7700
7705
7710
7715
7720
7725
7730
7735
7740
7745
7750
7755
7760
7765
7770
7775
7780
7785
7790
7795
7800
7805
7810
7815
7820
7825
7830
7835
7840
7845
7850
7855
7860
7865
7870
7875
7880
7885
7890
7895
7900
7905
7910
7915
7920
7925
7930
7935
7940
7945
7950
7955
7960
7965
7970
7975
7980
7985
7990
7995
8000
8005
8010
8015
8020
8025
8030
8035
8040
8045
8050
8055
8060
8065
8070
8075
8080
8085
8090
8095
8100
8105
8110
8115
8120
8125
8130
8135
8140
8145
8150
8155
8160
8165
8170
8175
8180
8185
8190
8195
8200
8205
8210
8215
8220
8225
8230
8235
8240
8245
8250
8255
8260
8265
8270
8275
8280
8285
8290
8295
8300
8305
8310
8315
8320
8325
8330
8335
8340
8345
8350
8355
8360
8365
8370
8375
8380
8385
8390
8395
8400
8405
8410
8415
8420
8425
8430
8435
8440
8445
8450
8455
8460
8465
8470
8475
8480
8485
8490
8495
8500
8505
8510
8515
8520
8525
8530
8535
8540
8545
8550
8555
8560
8565
8570
8575
8580
8585
8590
8595
8600
8605
8610
8615
8620
8625
8630
8635
8640
8645
8650
8655
8660
8665
8670
8675
8680
8685
8690
8695
8700
8705
8710
8715
8720
8725
8730
8735
8740
8745
8750
8755
8760
8765
8770
8775
8780
8785
8790
8795
8800
8805
8810
8815
8820
8825
8830
8835
8840
8845
8850
8855
8860
8865
8870
8875
8880
8885
8890
8895
8900
8905
8910
8915
8920
8925
8930
8935
8940
8945
8950
8955
8960
8965
8970
8975
8980
8985
8990
8995
9000
9005
9010
9015
9020
9025
9030
9035
9040
9045
9050
9055
9060
9065
9070
9075
9080
9085
9090
9095
9100
9105
9110
9115
9120
9125
9130
9135
9140
9145
9150
9155
9160
9165
9170
9175
9180
9185
9190
9195
9200
9205
9210
9215
9220
9225
9230
9235
9240
9245
9250
9255
9260
9265
9270
9275
9280
9285
9290
9295
9300
9305
9310
9315
9320
9325
9330
9335
9340
9345
9350
9355
9360
9365
9370
9375
9380
9385
9390
9395
9400
9405
9410
9415
9420
9425
9430
9435
9440
9445
9450
9455
9460
9465
9470
9475
9480
9485
9490
9495
9500
9505
9510
9515
9520
9525
9530
9535
9540
9545
9550
9555
9560
9565
9570
9575
9580
9585
9590
9595
9600
9605
9610
9615
9620
9625
9630
9635
9640
9645
9650
9655
9660
9665
9670
9675
9680
9685
9690
9695
9700
9705
9710
9715
9720
9725
9730
9735
9740
9745
9750
9755
9760
9765
9770
9775
9780
9785
9790
9795
9800
9805
9810
9815
9820
9825
9830
9835
9840
9845
9850
9855
9860
9865
9870
9875
9880
9885
9890
9895
9900
9905
9910
9915
9920
9925
9930
9935
9940
9945
9950
9955
9960
9965
9970
9975
9980
9985
9990
9995
10000
10005
10010
10015
10020
10025
10030
10035
10040
10045
10050
10055
10060
10065
10070
10075
10080
10085
10090
10095
10100
10105
10110
10115
10120
10125
10130
10135
10140
10145
10150
10155
10160
10165
10170
10175
10180
10185
10190
10195
10200
10205
10210
10215
10220

into audible voice that can be communicated to a user. (Alternatively, the platform might output the digitized speech signal itself, such as in a voice-over-IP communication system).

5 A voice command application may thus specify text that represents voice prompts to be spoken to a user. When the voice command platform encounters an instruction to speak such text, the platform may provide the text to the TTS engine. The TTS engine may then convert the text to an outgoing digitized speech signal, and the platform may convert the signal to analog speech and send it to the user. In converting from text to speech, the TTS engine may also make use of the dictionary database of phonemes, so that it can piece together the words that make up the designated speech.

Also on the outgoing side, a voice command platform may include a set of stored voice prompts, in the form of digitized audio files (e.g., *.wav files) for instance. These stored voice prompts would often be common prompts, such as "Hello", "Ready", "Please select from the following options", or the like. Each stored voice prompt might have an associated label (e.g., a filename under which the prompt is stored). And, by reference to the label, a voice command application might specify that the voice command platform should play the prompt to a user. In response, the voice command platform may retrieve the audio file, convert it to an analog waveform, and send the analog waveform to the user.

20 A voice command application can reside permanently on the voice command platform, or it can be loaded dynamically into the platform. For instance, the platform can include or be coupled with a network or storage medium that maintains various voice command applications. When a user calls the platform, the platform can thus load an application from the storage medium and execute the application. Further, in response to logic in the application (such as logic keyed to a user's response to a menu of options), the platform can load and execute another

application. In this way, a user can navigate through a series of applications and menus in the various applications, during a given session with the platform.

5 A voice command application can be written or rendered in any of a variety of computer languages. One such language is VoiceXML (or simply "VXML"), which is a tag-based language similar the HTML language that underlies most Internet web pages. (Other analogous languages, such as SpeechML and VoxML for instance, are available as well.) By coding a voice command application in VXML, the application can thus be made to readily access and provide web content, just as an HTML-based application can do. Further, when executed by the voice command platform, the VXML application can effectively communicate with a user through speech.

10 An application developer can write a voice command application in VXML. Alternatively, an application developer can write an application in another language (such as Java, C, C++, etc.), and the content of that application can be rendered in VXML. (For instance, when the platform loads an application, the platform or some intermediate entity could transcode the application from its native code to VXML.)

15 In order for a voice command platform to execute a VXML application or other tag-based application, the platform should include a VXML browser or "interpreter." The VXML interpreter functions to interpret tags set forth in the application and to cause a processor to execute associated logic set forth in the application.

20 A VXML application can be made up of a number of VXML documents, just like an HTML web site can made up of a number of HTML pages. A VXML application that is made up of more than one document should include a root document, somewhat analogous to an HTML home page. According to VXML, the root document defines variables that are available

document of the application, the voice command platform should responsively execute the logic that accompanies that grammar in the root document of the application.

In a voice command platform, each navigation point may have a respective identifier or label. For example, each voice command application can have a respective label, such as a network address where the application is maintained. As another example, a voice command application can define a number of successive menus through which a user can browse, and each menu might have a respective label by which it can be referenced. A voice command platform can use these labels to move from application to application or from menu item to menu item, just as hyperlinks operate to cause a browser to move from one web page (or component of one web page) to another.

In VXML, for instance, each VXML document will have a respective Universal Resource Identifier (URI), which is akin to a Universal Resource Locator (URL) used to identify the network location of an HTML page. A given VXML document may thus define logic that instructs the voice command platform to load and execute another VXML document from a designated URI. For instance, a VXML document may indicate that, if a user speaks a particular grammar, the platform should load and execute a particular VXML document from a designated URI, but that, if the user speaks another grammar, the platform should load and execute another VXML document from another designated URI.

An example of a VXML application is a weather reporting application. The weather reporting application may have a root document that includes a tag defining a welcome message and prompting a user to indicate a city or zip code. The root document may further set forth a bundle of grammars that are possible city names and corresponding zip codes that a user can speak in response to the prompt.

When the voice command platform executes this root document, it may thus send the welcome message/prompt to the TTS engine to have the message/prompt spoken to the user. In turn, when the user speaks a response, the SR engine would identify the response as one of the acceptable grammars. The platform would then continue executing the root document in view of the spoken response.

The root document might next prompt the user to indicate whether the user would like to hear today's weather or an extended forecast, and the user would again speak a response. In turn, the root document might indicate that, if the user selects "today's weather," the platform should load and execute a subsidiary document from a designated URI, and if the user selects "extended forecast," the platform should load and execute a different subsidiary document from another designated URI. Of course, many other examples of VXML applications are possible as well.

In most cases, a platform provider will own and operate the voice command platform. Content providers (or independent application developers having a relationship with a content provider for instance) will then provide the VXML applications to be executed by the platform. The platform provider may also provide some applications for the platform and may therefore function as a content provider as well.

Further, a content provider or other application developer can personalize a VXML application, through reference to user profiles. For example, a telecommunications service provider (e.g., local exchange carrier or interexchange carrier) can provide a voice-activated-dialing (VAD) application that allows users to dial a telephone number by speaking a name. To support this feature, the VAD application may direct the voice command platform to prompt a user for a user ID or to determine the user ID based on calling number identification provided when the user's call was connected to the platform. The VAD application may then instruct the

platform to call up a personalized VAD application (through use of Microsoft Active Server Pages, for instance), which is tied to the user's personal address book. Each name in the address book may then define an acceptable grammar. When the user speaks one of the names, the application may cause the platform to retrieve a corresponding telephone number and to provide that number to a network switch to facilitate initiating the call.

SUMMARY

The present invention is directed to a method and system for consolidated message notification in a voice command platform. Generally speaking, the platform maintains, for each user, a summary of counts of messages waiting for the user at multiple message portals. The platform receives update messages from the message portals, indicating updates of the counts of messages waiting, and the platform revises its summary accordingly. During a voice command session with the user, such as at the initiation of the session, the platform then presents an indication of the summary of message counts to the user, such as by sending to the user a speech signal indicating the summary.

Thus, in one respect, the invention may take the form of a voice command platform that is programmed to receive a call from a user, to answer the call, and to send to the user a speech signal representing a consolidated summary of counts of messages waiting for the user at multiple message portals (e.g., e-mail portals, voice mail portals, fax portals, and/or instant messaging portal).

Further, the platform may be programmed to maintain a record of counts of messages waiting for the user at the message portals. In this regard, for instance, the platform may be programmed to receive from the message portals update signals that indicate updates of counts of

messages waiting for the user at the message portals. (The portals may send the update signals autonomously, and/or the platform may request the updates, such as with periodic polling of the portals.) The platform may then programmatically update the record based on the update-signals.

5 In addition, the platform may be programmed to personalize the consolidated summary, based on a notification profile maintained for the user (e.g., in a user profile record accessible by the platform). The notification profile might indicate parameters to include in the consolidated summary, and the platform then includes those parameters in the consolidated summary. Further, the notification profile might indicate user-specific names of the message portals, and the consolidated summary may then refer to the portals by those names.

10 In another respect, the invention may take the form of a voice command platform that includes a user communication interface, a processor, an application-processing module, a voice-processing module, a user profile store, and consolidated-message-notification logic. The application-processing module is executable by the processor to process voice command applications that define user-prompts, allowed grammars, and application logic. The voice-processing module is executable by the processor to recognize the allowed grammars in speech signals received from a user via the user communication interface, and to convert the user-prompts into speech signals for transmission to the user via the user communication interface.

20 The user profile store includes, respectively for each user, a consolidated message summary indicating counts of messages waiting for the user at a plurality of message portals. In turn, the consolidated-message-notification logic is executable by the processor to communicate to a given user (e.g., at the initiation of a voice command session with the user – such as in a welcome message, or just after a welcome message, spoken to the user), via the user

communication interface, an indication of the counts of messages waiting for the user at the plurality of message portals, as reflected by the message summary for the given user.

In this arrangement, the processor may receive update-messages from multiple message portals. The update messages may indicate updated counts of messages waiting at the message
5 portals for at least one user. Based on the update-messages, the processor may then update the consolidated message summary for at least one user.

Further, the platform may include portal-polling logic executable by the processor to request (e.g., periodically) updates of message waiting counts from the plurality of message portals. And the processor may then receive the update-messages in response.

In yet another respect, the invention may take the form of a method that involves receiving from each of multiple separate message-portals a respective message-waiting count for a common user, receiving a voice call from the common user, and presenting to the common user a spoken summary of the respective message-waiting counts for the plurality of separate message-portals. Further, the method may involve polling the message-portals for respective message-waiting counts for the common user.

The process of receiving a message-waiting count from a given message-portal may thus involve receiving a signaling message that specifies the respective message-waiting count. Alternatively, the process may involve receiving multiple signaling messages that cooperatively define the respective message count. (E.g., each message may be an update of a prior count, or
20 may add to or modify a prior count.)

When speaking the consolidated summary to the common user, the method may also involve reading from a stored notification profile for the common user, to identify a user-specific

name for at least a given one of the message-portals. In turn, the method may involve referring to that portal by that name, in the spoken summary.

These as well as other aspects and advantages of the present invention will become apparent to those of ordinary skill in the art by reading the following detailed description, with appropriate reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Exemplary embodiments of the present invention are described herein with reference to the drawings, in which:

Figure 1 is a functional block diagram illustrating the layers of a system in which the exemplary embodiments can be employed; and

Figure 2 is a functional block diagram of a voice command platform including enhanced system logic in accordance with the exemplary embodiments.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

1. Exemplary Voice Command System

Referring to the drawings, Figure 1 is a functional block diagram illustrating the layers of a system in which an exemplary embodiment of the present invention can be employed. As shown in Figure 1, the system may be considered to include three layers, (i) a user layer 12, (ii) a platform or system layer 14, and (iii) an application layer 16. The user layer 12 provides a mechanism 18, such as a telephone, for a person to communicate by voice with the platform. The system layer, in turn, includes a user communication interface 20, a processor 22 (i.e., one or

more processors), a voice-processing module 24, a VXML interpreter module 26, and cache 28. Application layer 16 then defines an application 30, which may be made up of a root document 32 and subsidiary documents 34 that can be loaded into cache 28 and executed by processor 22, i.e., by the voice command platform.

5 User communication interface 20 may take various forms. For example, the user communication interface can provide a circuit or packet interface with a telecommunications network (such as the PTSN or the Internet, either of which may provide connectivity to one or more users via at least one wireless communications link – such as a cellular communication system for instance). The communication interface may, in turn, include an A-D and D-A converter (not shown) as described above, for converting between analog signals on the user side and digital signals on the platform side. Processor 22 then sends and receives communications via user communication interface 20.

 Voice-processing module 24 and VXML interpreter module 26 preferably define program instructions that can be executed by processor 22 and data that can be referenced by the processor, to carry out basic voice platform functions. All of this program logic can be stored in suitable data storage, such as ROM or a disk drive for example.

 For instance, voice-processing module 24 preferably defines an SR engine 36 and a TTS engine 38, as well as a voice prompt store 40. Further, voice-processing module 24 may include a phoneme dictionary 42 that the processor can reference to facilitate operation of the SR engine and the TTS engine. VXML interpreter module 26, in turn, may facilitate execution of application 30, such as by interpreting tags in the application.

 Cache 28 functions to temporarily hold application 30 (or one or more documents of the application) when processor 22 is executing the application. Thus, in operation, processor 22

may retrieve application 30 from a designated URI on the Internet (or elsewhere) and may load the application into cache 28. The processor may then execute the application, using VXML interpreter 26 to interpret tags, using TTS engine 38 and voice prompt store 40 to send speech to a user, and using SR engine 36 to recognize speech spoken by a user.

5 It should be understood that that this and other arrangements described herein are set forth for purposes of example only. As such, those skilled in the art will appreciate that other arrangements and other elements (e.g., machines, interfaces, functions, orders and groupings of functions, etc.) can be used instead, and some elements may be omitted altogether. Further, as in most telecommunications applications, those skilled in the art will appreciate that many of the elements described herein are functional entities that may be implemented as discrete or distributed components or in conjunction with other components, and in any suitable combination and location.

For example, although the voice command system is shown to include a VXML interpreter, the system may include another type of voice browser. Alternatively, the system might not be browser-based at all. In that event, some of the functions described herein could be applied by analogy in another manner, such as through more conventional interactive voice response (IVR) processing. Other examples are possible as well.

2. **Enhanced System Layer**

20 In accordance with the exemplary embodiment, system layer 14 can be enhanced in several respects, so as to provide a more robust voice command platform. Figure 2 depicts an expanded functional block diagram of the system layer 14, showing some of the improvements contemplated.

As shown by way of example in Figure 2, exemplary system layer 14 includes the components described above, such as a user communication interface 20, a processor 22, voice-processing 24, VXML interpreter 26, and cache 28. In addition, according to the exemplary embodiment, the system layer includes or has access to an enhanced set of system logic, preferably comprised of software modules that can be executed by the processor and data that can be referenced by the processor.

Some or all of the enhanced system logic can be co-located with the platform or can be located elsewhere, such as on an accessible network. Further, aspects of this enhanced system logic may be integrated together with the voice-processing module 24 and/or with the VXML interpreter module 26 to any extent desired.

Referring in more detail to Figure 2, the enhanced system logic can include multiple text-to-speech engines ($TTS_1, TTS_2, \dots TTS_N$) 44 (e.g., separate sub-routines of a main TTS module, or separate TTS modules altogether) that may be selectively applied by processor 22 (i.e., selected by the processor in response to a designated stimulus or instruction). The TTS engines can differ from each other in terms of language (e.g., English, Spanish, etc.), dialect (e.g., accent), voice (e.g., male, female) and/or other aspects, so as to speak to users in different manners. (In the exemplary embodiment, the platform can similarly include multiple speech recognition engines (not shown), or multiple core phoneme dictionaries, to facilitate converting speech of various types to data.)

Additionally, the enhanced system logic can include multiple voice prompt stores or sets of voice prompts 46 that can be selectively applied by the processor. Each prompt set may differ from the other prompt sets in various ways, such as being spoken in different dialects or different

voices in a common language, or being spoken in different languages, or with different background music, for instance.

For example, one prompt set might comprise the prompts spoken in a particular celebrity voice (e.g., in a recognizable voice of a particular movie or television star), and another prompt set might comprise the prompts spoken in a different celebrity voice. Still another prompt set might comprise the prompts spoken in a male voice, and another might comprise the prompts spoken in a female voice. Further, another might comprise the prompts spoken in a high energy voice, while another might comprise the prompts spoken in a low key voice. And still another might comprise the prompts spoken in English, while another might comprise the prompts spoken in French.

On the platform, the prompt sets can be stored in separate directories (folders) on a disk drive, where each directory has a respective name to designate the prompt set. For instance, the directories may be named "Celebrity1", "Celebrity2", "male", "female", "high energy", "low key," etc. This can facilitate quick and easy retrieval of desired prompts. Other mechanisms for storing and identifying the various prompt sets, however, are possible as well.

As further shown in Figure 2, the enhanced system logic can include a set of one or more secondary phoneme dictionaries 48. Each of these secondary dictionaries is preferably an application-specific (or content-provider-specific) phoneme dictionary, provided by a content provider to improve voice processing functions when the platform is executing the respective application.

Additionally, the enhanced system logic can include a set of system-level grammars 50. The system-level grammars 50 are preferably grammars that are defined for the voice command platform generally. As such, the voice command platform can recognize the system-level

grammars when spoken by a user at substantially any navigation point, i.e., while any voice command application is being executed by the platform.

In addition, the enhanced system logic can include a user profile store 52. For various aspects of the exemplary embodiment, the user profile store 52 may define (i) user ID information 54, (ii) user preference information 56, (iii) user session information 58, and (iv) user consolidated-messaging information 60. This information can be maintained in a relational database or in any other manner desired (whether as one module or separate objects, such as a meta-directory for instance), and it can be co-located with the rest of the platform or can be located elsewhere.

User ID information 54 may include any information that the platform can use to identify a user, so as to facilitate providing personalized voice command services across multiple applications. In the exemplary embodiment, user ID information will be ANI information, such as the billing number or telephone number of the calling party. Alternatively, however, user ID information can also take other forms, such as username and password.

User preference information 56 can also take various forms. For example, user preference information may include a set of personal grammars/aliases, which define grammars specific to a user. As another example, user preference information may include an expert mode flag (or flags) for the user, which indicates whether or not the platform should interface with the user according to an "expert mode." As still another example, user preference information may include indications of preferred content providers, such as a preferred e-mail provider, a preferred voice-mail provider, a preferred weather-reporting provider, etc.

As yet another example, user preference information may include a set of bookmarks/favorites (with associated names), which function as user-specific shortcuts to URIs

or menu items. And as still a further example, user preference information may include a user-specific speech persona, which indicates a voice or persona for the voice-processor to use in communicating to the user (e.g., an indication of which of several possible pre-recorded versions of a prompt to play to the user and/or which TTS engine to use for the user).

5 User session information 58 can also take various forms. For example, user session information can comprise a record of a user's current session state. In this regard, session state information may include an indication of whether or not the voice-processor (or, more specifically, the speech recognition engine) is currently muted for the user's session. And as still another example, session state information may include an indication of which of the multiple TTS engines 44 and/or voice prompt sets 46 is currently active (selected, or being used) for the user or for the user's session.

 User session information 58 can also comprise a record of a user's navigation history. For example, the navigation history record may specify in order the last predetermined number of navigation points (e.g., URIs, menu items, etc.) that the user has accessed, preferably including the latest navigation point accessed by the user. Further, the navigation history record may include a summary ("rollup") indication of the number of times that the user has accessed each navigation point. This navigation history information can be used as a basis to determine whether the system should automatically switch the user to expert mode, as well as to restore a user's voice command session in the event a system disconnect occurs.

20 User consolidated-messaging information 60, similarly, can take various forms. In the exemplary embodiment, the consolidated-messaging information comprises a summary the number of messages waiting for the user at multiple separate message portals. For instance, the consolidated-messaging information for a user can indicate a count of voice mail messages

waiting for the user at one message portal, a count of e-mail messages waiting for the user at another message portal, and account of fax messages waiting for the user at another message portal. The user profile may also identify each of the user's message accounts (e.g., POP3 e-mail account information, etc.)

5 As still further shown in Figure 2, enhanced system logic preferably defines a number of additional logic modules 62 that are executable by the processor to carry out enhanced functions described herein. These enhanced functions can involve, for instance, (a) providing enhanced response to voice commands, (b) allowing dynamic control over the voice-processing module, (c) monitoring and use of navigation history, and (d) providing unified messaging notification.

10 Some of these functions may involve having processor 22 refer to user profile store 52 in order to determine a user preferences or other user-specific information. In order to facilitate this, a mechanism should preferably be provided to identify a user who contacts the platform. This mechanism can take various forms.

15 As noted above, for example, user profile store 52 preferably indicates as user ID information each user's ANI. When a user places a conventional telephone call to the voice command platform, the processor can receive the user's ANI in the call setup signaling or otherwise. The processor may then query the user profile store 52 to find the user's profile record. In turn, the processor can store in memory a pointer to the user's profile record. Alternatively, the processor can extract parameters from the user's profile record and store those
20 parameters in memory for quicker access.

 As another example, as also noted above, user profile store 52 could indicate a username and password for each user. In an alternative arrangement, the processor could be programmed to prompt a user to specify a username when the user calls the platform, and the processor can

then query the user profile store to find a matching record. Further, the processor can prompt the user to specify a password, and the processor can then authenticate the user by reference to the stored password in the profile record.

This alternative identification function can be carried out by having the processor execute a system-level VXML application when a caller calls the platform. The system-level application may define a username voice prompt such as "Welcome to the voice command system. Please state your name to begin." In turn, the system-level application can receive a response spoken by the user and can direct the processor to search the user profile for a matching record.

Each of the exemplary enhanced functions will now be described in the following sections.

a. Providing Enhanced Response to Voice Commands

In accordance with the exemplary embodiment, the enhanced system logic provides for enhanced response to voice commands. Generally speaking, the platform receives a voice command during a voice command session with a user, and the platform then responsively carries out a set of system-level logic.

As examples, enhanced response to voice commands can involve (i) recognizing system-level grammars (i.e., system-specific grammars) that can be spoken by a user, (ii) recognizing personal grammars (i.e., user-specific aliases of other grammars) that can be spoken by a user, and (iii) hosting a set of personal bookmarks (i.e., user-specific shortcuts to navigation points). These examples will be described in the following sections.

i. Responding to User Utterance of System-Level Grammars

Normally, an *application* that is executed by the voice command platform can define a set of grammars that the system layer will recognize. For instance, a weather reporting

application might define "today's weather" and "extended forecast" as grammars that a user can speak, and the application might then define logic to be executed respectively in response to each of those grammars.

In accordance with the exemplary embodiment, as noted above, enhanced system logic
5 40 can include a set of system layer grammars, which are grammars that the system layer (e.g., VXML interpreter) will recognize globally, i.e., substantially regardless of what application the system is currently executing.

The system-level grammars, and their associated logic, can take various forms. For example, the processor can be programmed to recognize as system-level grammars the terms "Voice Command" and "Main Menu" and to respond to either of these grammars by presenting the user with the initial ready prompt, i.e., the prompt that the user might first encounter when the user calls the platform. As another example, the processor can be programmed to recognize as a system-level grammar the word "Mute" and to respond to this grammar by turning off (or partially turning off) the speech recognition engine 36. Similarly, the processor can be programmed to recognize the word "Unmute" and to respond to this grammar by turning on the speech recognition engine 36.

As still another example, the processor can be programmed to recognize the grammars "Save" and "Bookmark" and to respond to either of these grammars by saving in the user's profile a bookmark to the current navigation point, or by executing a subsidiary bookmark-
20 application. (As will be noted below, the processor may prompt the user to specify a shortcut name for the bookmark, which the processor may also store.) And as yet another example, the processor may be programmed to recognize the grammar "Voice Command Help" and to respond to this grammar by executing a voice command help module/application. Still another

example may be a "Pause" or "Go to Sleep" system-level grammar, which may direct the processor to pause the current session.

In the exemplary embodiment, the system-level grammars and their associated logic can be defined in a system-level application that gives the system-level grammars a global scope for subsidiary applications. For example, the SR engine can be hard-coded to recognize the set of system-level grammars. Alternatively, in a VXML-based platform for instance, the system-level application can be a root VXML application that the platform is programmed to run by default when a user calls the platform (preferably without having to be referenced by another VXML document).

The root VXML application can function to welcome the user to the platform, such as by providing an initial ready prompt to the user and providing an initial menu of choices (such as VAD, news, sports, stocks, etc.) and calls to subsidiary applications. And the root VXML application can ideally define the system-level grammars and their associated logic. Thus, as subsidiary applications are called from the root document and from subsidiary documents, the system-level grammars and associated logic defined in the root VXML application will remain available for use in all subsidiary documents.

Preferably, if a content provider defines a grammar in an application that conflicts with (e.g., is the same as) a system-level grammar, the system-level grammar will take precedence. I.e., platform would be programmed to execute the logic defined for the system-level grammar, rather than the logic defined for the conflicting application-level grammar. However, exceptions could be possible.

As another example, the system layer can be programmed with one or more system-level grammars that function as "aliases," i.e., nicknames or shortcuts (or, for that matter, more

extensive names) for other grammars. For instance, the system layer might include in memory a table that translates between alias system-level grammars and actual grammars. When a user speaks one of the alias system-level grammars, the system layer can find a match in the table and responsively convert the spoken grammar into the associated actual grammar. The system-layer
5 can then treat the actual grammar as having been spoken by the user.

Advantageously, because this aliasing mechanism exists in the system layer, the voice command platform can provide translations between aliases and actual grammars globally for all applications (or, allowing for exceptions, substantially all applications) executed by the platform. For example, the system-layer might define the terms "weather outlook" and "tomorrow's weather" as alias grammars for the word "extended forecast." In turn, when the system is executing a weather-reporting application and a user speaks "weather outlook," the system may convert "weather outlook" into "extended forecast" and treat "extended forecast" as being spoken by the user. The system may thus execute the logic that the application defines for the response "extended forecast," such as calling up an extended-forecast subsidiary URI, for instance.

ii. Responding to User Utterance of Personal-Grammars

In the exemplary embodiment, as noted above, user profile store 52 can also define a set of user-specific grammars, i.e., personal grammars. In the exemplary embodiment, these personal grammars function as aliases for other grammars, similar to the system-level alias
20 grammars described above.

Thus, a user's profile preferably includes a table or other mechanism that translates between personal grammars and actual grammars. When a user speaks a personal grammar, the system can then find a match in the table and responsively convert the spoken grammar into the

associated actual grammar. The system layer can then treat the actual grammar as having been spoken by the user.

As an example, an application may define "stocks" as a grammar that a user can select to go to a stock-reporting application. Advantageously, the user's profile may define the personal grammar "check portfolio" to correspond to the grammar "stocks." When the user speaks "check portfolio," the system can thus translate the grammar into "stocks" and proceed to execute the application as if the user has spoken "stocks." Other examples are possible as well.

iii. Centrally Maintaining Personal Bookmarks

In accordance with the exemplary embodiment, as noted above, user profile store 52 can further include, for each user, a set of one or more bookmarks that function as user-specific shortcuts to navigation points, such as URIs or menu items. Advantageously, by maintaining these bookmarks centrally on the voice command platform (or otherwise in a manner accessible by the platform), a user can make use of the bookmarks when the user calls the platform from any location (absent other restrictions).

Preferably, the enhanced system logic provides a mechanism for maintenance of these bookmarks, such as recording the bookmarks, receiving user requests to use the bookmarks, and applying the bookmarks to navigate to a designated navigation point. In the exemplary embodiment, this can be accomplished through a system-level application (e.g., a VXML application) that a user can access by speaking a system-level grammar such as "Bookmark." I.e., the platform can be arranged to execute a bookmark-management application in response to the system-level grammar "Bookmark."

0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99

An exemplary bookmark-management application can provide a user with options such as (i) Save the current URI, (ii) Recall a bookmark, and (iii) Edit a bookmark. In response to user selection of the Save option, the bookmark-management application may direct the processor to save in the user's bookmark list the URI of the VXML document that the user was accessing before the user called the bookmark-management application (or the label of whatever other navigation point the user was accessing at the time). Further, the bookmark-management application may direct the processor to prompt the user for a shortcut name (i.e., an alias) for the bookmark, possibly providing the user with a suggested default name. Once the user speaks the name or agrees to use the default, the processor may store the name in the user's bookmark list together with the URI.

In response to user selection of the Recall option, the bookmark-management application may direct the processor to prompt the user to speak the name of a bookmark on the user's list. Once the user speaks the name, the processor may refer to the bookmark list to determine the corresponding URI and may then load the document from that URI and execute the document.

As a further enhancement, the bookmark-management application may direct the processor to recognize the command "Recall" followed by the name of the bookmark that the user wants to recall, and to responsively load and execute the corresponding URI. For instance, in response to a user speaking "Recall Shopping," the bookmark-management application may direct the processor to refer to the user's bookmark list to find a URI that corresponds to the name "Shopping." The processor may then find the corresponding URI and load and execute the URI. (Note that the "Recall URI" grammar structure can itself be made a system-level grammar of the voice command platform, so that a user can call up a stored URI at any time, without first expressly calling up the bookmark-application.)

In response to user selection of the Edit option, the bookmark-management application may direct the processor to prompt the user to speak the name of a bookmark on the user's list. The processor may then prompt the user to indicate whether the user wants to rename or delete the entry. Other examples are possible as well.

5 **b. Allowing Dynamic Control over
the Voice-Processing Module**

In accordance with the exemplary embodiment, the enhanced system logic also provides for dynamic control over the voice-processing module. Generally speaking, according to this embodiment, the voice-processing module 24 may have a number of selectable modes of operation. During a voice command session with a user, the platform may then receive a specification that corresponds with a mode of operation of the voice-processing module. In response to the specification, the processor 22 will cause the voice-processing module 24 to operate according to the mode of operation corresponding to the specification.

Advantageously, the platform may receive the specification from various sources. For example, the platform may receive the specification from a user during the voice command session (e.g., the user can speak the specification or enter the specification through DTMF key entry). As another example, the platform may receive the specification from a voice command application being processed during the voice command session (e.g., the application can command the platform to switch to a particular voice-processing mode). As still another
20 example, the platform can retrieve the specification from user profile store 52 (e.g., at the initiation of the voice command session (where "during" could be just before the platform answers a call from a user, for instance)).

Further, the particular voice-processing mode or change of voice-processing mode may also take various forms. For instance, in response to the specification, the platform might (i) mute the SR engine 36, (ii) selectively apply a designated text-to-speech engine, (iii) selectively apply a particular voice prompt store, (iv) selectively apply a particular speech persona, and/or
5 (v) selectively apply a designated secondary phoneme dictionary. Examples of these enhancements will be described in the following sections.

i. Muting the SR Engine

In an exemplary voice command platform, processor 22 is programmed to be able to selectively turn on or off the SR engine 36. When the SR engine is turned off, the system will not recognize words spoken by the user or will recognize only very specific grammars (or a limited set of grammars). For instance, the SR engine may be set to simply not analyze incoming digitized speech waveforms, or the SR engine may be set to analyze only the digitized speech waveforms in search of only very specific grammars.

Advantageously, enhanced system logic 40 enables a user to control the state of the speech recognition engine during the user's session with the voice command platform. As described above, for instance, a system-level grammar such as "Mute" can be defined. When the user says "Mute," the processor may respond by setting a flag in memory to indicate that the SR engine is turned off. The SR engine module 36 may then include logic to indicate that, if the flag indicates the SR is turned off, the processor should not apply the SR engine module or should
20 apply only designated aspects of the module.

In the exemplary embodiment, the user can then turn the speech recognition engine back on for the user's session, by entering a designated DTMF tone that the processor is programmed to recognize. Alternatively, when the speech recognition engine is off, the processor can be

programmed to recognize only the word "Unmute" as a command to cause the processor to turn back on the SR engine. For instance, the word "Unmute" can be a system-level grammar. Further, when the SR engine is off, it can still function to await the word "Unmute."

With this enhancement, a user can thus disable the speech recognition engine whenever the user wants to do so. For instance, if a user is listening to a long news story and is standing in a noisy airport, the user can say "Mute" to prevent a "barge-in" error such as where the background noise is erroneously recognized by the speech recognition engine. Advantageously, the word "unmute" is uncommon enough that it is not likely to be spoken unless a user intends to turn back on the SR engine.

**ii. Using an Application-Specific
Phoneme Dictionary**

As noted above, the system layer traditionally includes or has access to a core phoneme dictionary 42, which defines phonemes that the voice-processor 24 can use in performing speech recognition. This dictionary thus enables the voice command system to recognize words spoken by a user.

A problem with this arrangement, however, is that the core dictionary might not define all of the possible utterances that users may make in connection with applications executed by the platform. For instance, a mapping application might include as grammars a list of street names and city names, many of which might have unconventional pronunciations. In order for the voice command platform to be able to interface with a user, the platform should ideally be able to understand utterances that are specific pronunciations unique to the content provider or application. At the same time, a platform provider should not have to overhaul its core dictionary in order to be compatible with all possible applications.

To overcome this problem, as noted above, then enhanced system logic can include one or more secondary phoneme dictionaries 48, each of which can be associated with a particular application (e.g., a particular URI or other navigation point) and/or a particular content provider. In the exemplary embodiment, each secondary phoneme dictionary can be a uniquely named data file that defines additions to the core dictionary, such as additional phonemes that the processor should recognize when executing the SR engine 36. For instance, the secondary dictionary associated with one application might be named "Phoneme_46", while the secondary dictionary associated with another application might be named "Phoneme_47".

Thus, when the processor executes an application, the processor can determine whether the application has an associated secondary dictionary 48. If so, the processor can apply that secondary dictionary 48 in conjunction with the core dictionary 42 as the processor executes the SR engine 36.

To associate each secondary phoneme dictionary with an application, the voice command platform can include or have access to a translation table or other logic that indicates, for a given application, which secondary phoneme dictionary to use. The table or other logic can be set up to correlate a given application with a given secondary phoneme dictionary and/or to correlate all applications by a given content-provider with a given secondary phoneme dictionary. Further, the table can correlate more than one secondary dictionary with a given application. Thus, to determine whether an application has an associated secondary phoneme dictionary, the processor can refer to the translation table or other logic.

The voice command platform can obtain secondary phoneme dictionaries in various ways. In the exemplary embodiment, for instance, a content-provider or other entity can provide

the voice command platform with a secondary phoneme dictionary to be used for a given application (or applications).

For example, the content provider can provide the secondary phoneme dictionary in advance, i.e., before the platform loads and executes the application. Alternatively, an application can include a secondary phoneme dictionary as a component that the platform will load when the platform loads the application. (This can be similar to how a browser loads a graphic or other component when it loads a conventional HTML web page.) The enhanced system logic will then cause the processor to load the secondary dictionary, store the secondary dictionary for use during execution of the application, and use the secondary dictionary. In the latter scenario, the processor can be programmed to delete the secondary dictionary when the processor finishes executing the application, or the processor can maintain the secondary dictionary for use again next time the application is executed.

iii. Applying User-Specific Designation of Speech Persona

According to the exemplary embodiment, as noted above, user profile store 52 can indicate for each user a user-specific persona. A user-specific persona defines a voice façade by which the platform will present itself to the user. As such, the persona could define a specific TTS engine (of TTS engines 44) and/or voice prompt store (of voice prompt stores 46) that the platform should apply when it speaks to the user. Other aspects of a user-specific speech persona are possible as well. For instance, a given user's speech persona may dictate a particular tone and/or pitch for the TTS engine to use when generating analog speech signals.

The persona is user-specific, in that the platform could use one persona when interacting with one user and a different persona when interacting with another user. In the exemplary

embodiment, the user profile record for a given user can include a table or other indication that identifies a TTS engine and/or voice prompt set to use during sessions with the user and that may further indicate other parameters about the way that the speech should be presented to the user. Further, the persona designation can define more complex persona-selection logic, such as that the platform should use one persona when interacting with the user at one time (e.g., daytime) and a different persona when interacting with the user at another time (e.g., evening).

The persona designation for a given user may point specifically to a TTS engine and/or voice prompt set, or it may point more generally to a persona that the platform should use. For example, the persona designation may indicate the name of a particular TTS engine (e.g., as "English_tts", "Spanish_tts", "male_tts", "female_tts", etc.) and/or the label of a particular voice prompt directory (e.g., "Celebrity_1", "Celebrity_2", "male", "female"). Alternatively, the persona designation may indicate more generally a persona category, in which case the processor may refer to a stored translation table that indicates which TTS engine and/or voice prompt set to use so as to achieve the designated persona.

Further, the voice of the prompt set and the voice of the TTS engine can be made to approximate or match each other (e.g., both can be English speaking male voices of roughly the same pitch), so that a user profile record can indicate a TTS engine and the platform can select a corresponding voice prompt set (and vice versa). To do so, the platform can be programmed with system-level logic (e.g., a cross-reference table) that matches voice prompt sets with TTS engines and that indicates parameters such as pitch.

Thus, when a user calls the platform, processor 22 can programmatically consult user profile store 44 to determine which persona is to be used for speaking to the user. As the processor then executes an application for the user, the processor may then apply the designated

TTS engine and/or voice prompt set. Further, the platform can have a default persona (e.g., a standard "male" persona), which the processor can be programmed to use when the user's profile does not indicate that the platform should use another persona.

In the exemplary embodiment, a platform can also allow a user to change the user's persona selection mid-session. For instance, by speaking a system-level grammar such as "Change Persona", the user could direct the platform to execute a persona-management application. The persona-management application can tell the user which persona is currently set for the user and can prompt the user to select a different persona. In response to user-selection of a different persona, the processor can update the user's profile record accordingly and can begin to use the different persona.

Advantageously, a user-specific speech persona can be tailored to separately suit each particular user. For instance, a female user might prefer a female voice TTS engine and a corresponding female voice prompt set, while a male user might prefer a male voice TTS engine and corresponding mail voice prompt set. Other examples are possible as well.

iv. Applying a Voice-Processing Mechanism Designated by an Application

In the exemplary embodiment, the enhanced system logic also enables an application to specify which of various voice-processing mechanisms the platform should use during execution of the application. For instance, the logic can enable an application to specify which of multiple TTS engines 44, voice prompt stores 46 and/or secondary phoneme dictionaries 48 to use during execution of the application.

This facilitate this, each voice-processing mechanism can have a unique identifier, and the processor can be programmed to recognize and respond to an instruction (e.g., a predefined

command, tag, etc.) in an application that specifies the use of a particular voice-processing mechanism. Each such instruction may, for instance, indicate a voice-processing type and a voice-processing value.

For example, as noted above, exemplary TTS engines might be named "English_tts", "Spanish_tts", "male_tts", "female_tts", etc., exemplary voice prompt stores might be named "Celebrity_1", "Celebrity_2", "male", "female", etc., and exemplary phoneme dictionaries might be named "Phoneme_46" and "Phoneme_47". Exemplary VXML tags to specify a voice-processing mechanism might then take the following form:

```
<TTS = "Spanish_tts">
<VOICEPROMPT = "Celebrity_1">
<PHONEME = "Phoneme_46">
```

An application can be written to include one or more of these tags. And the VXML interpreter can then be written to recognize these tags and to direct the processor to apply the designated voice-processing mechanism.

Further, the platform can include or have access to a translation table or other logic that correlates unique voice-processing keys (e.g., code numbers) with designated voice-processing mechanisms. According to that table, a given voice-processing key could be correlated with a specific TTS engine 44, a specific voice prompt store 46 and/or a specific secondary phoneme dictionary 48. For instance, voice-processing key "3752" can be correlated with a combination of the "male_tts" TTS engine and the "male" voice prompt store.

A suitable instruction in an application can then specify a particular voice-processing key, such as by a VXML tag like the following:

```
<VOICEPROC = "3752">
```

In turn, the VXML interpreter can be written to recognize such a tag and to direct the processor to apply the corresponding voice-processing mechanism(s).

With the exemplary arrangement, a content-provider can thus write a voice command application to take advantage of the ability to dynamically select a voice-processing mechanism that the platform should use. For instance, a given application may include a voice-processing tag at the start of its root document, which will direct the platform to apply a particular TTS engine, a particular voice prompt store and/or a particular secondary phoneme dictionary during execution of the application.

Further, by including a voice-processing tag in an application, the application can cause the platform to switch between voice-processing mechanisms dynamically during execution of the application. For instance, the application may cause the platform to switch from one TTS engine to another during execution of the application.

In the exemplary embodiment, the voice command platform will preferably be programmed to treat particular voice-processing mechanisms as default mechanisms, in the event no other designation is made. For example, the default setting may be to use the "English_tts" TTS engine, the "female" voice prompt store, and no secondary phoneme dictionary. Other examples are possible as well.

Finally, another way for an application to specify a voice-processing mechanism to use is for the application or content-provider to provide the platform with the voice-processing mechanism. As noted above, for example, the application or content-provider can provide the platform with a secondary phoneme dictionary 48. Further, since the TTS engines 44 and voice prompt stores 46 are also preferably software-based, it would be possible for an application or

content-provider to similarly provide the platform with an application-specific TTS engine and/or voice prompt store to use during execution of the application.

c. Monitoring and Use of Session Information

According to another aspect of the exemplary embodiment, the enhanced system logic also provides for maintaining and using information concerning user sessions with the voice command platform. In one respect, this functionality thus involves maintaining user session information. And in another respect, this functionality then involves using the session information to enhance user interaction with the platform in some manner. By way of example, two enhancements that a user's session information can facilitate are (1) automatically transitioning a user to expert mode based on the user's navigation history, and (2) automatically restoring a session that was abruptly cut off by a system disconnect or other event.

The process of maintaining user session information and using the information to perform these exemplary functions will now be described.

i. Maintaining User Session Information

As noted above, user profile store 52 can include a record of user session information, such as current session state information as well as user navigation history. In the exemplary embodiment, the processor 22 is preferably programmed to maintain this session information for each user, i.e., to keep track of each user's session with the platform.

In terms current session state information, for instance, the processor may be programmed to record in the user profile an indication of whether or not the speech recognition engine 36 is muted for a given user's session, as well as an indication of which TTS engine and/or voice prompt set is currently active for the user's session. These parameters can be set as flags or other parameters in the user's profile record.

In terms of navigation history, the processor may be programmed to maintain a list of the URIs accessed by the user, beginning with the most recent. In addition, as noted above, the processor may maintain a summary table that indicates, for each URI, how many times the user has accessed the URI. The processor may thus increment an entry in the summary table each time a user accesses the same URI.

Further, the processor may also be programmed to maintain in the user profile a record of how many times each user has contacted the voice command platform, possibly together with indications of the dates and/or times that the user has contacted the platform.

ii. Automatically Transitioning a User to Expert-Mode, Based on Navigation History

Generally speaking, a voice command application can define a number of voice prompts, acceptable response grammars, and associated logic that the platform is to execute in response to particular responses or other events. Thus, as a user interacts with the voice command application, the platform would present designated voice prompts to the user, receive responses from the user, and carry out logic keyed to the responses or other events.

A user who interacts regularly with the same application, however, may tire of listening to the same voice prompts over and over, and having to regularly provide the same response grammars to each voice prompt. To enhance the user experience, a robust voice command application can therefore define multiple sets of voice prompts, acceptable response grammars and/or associated logic.

As an example, one set might be considered a standard set, which would be appropriate for interaction with the typical user. And another set might be considered an expert set, which would be appropriate for interaction with a user considered to be an expert. Thus, depending on

whether the user who is interacting with the voice command application is designated as an expert user or not, the application may then instruct the processor to apply either the standard set or the expert set.

5 An expert set of voice prompts and/or acceptable response grammars may be more streamlined than the standard set, such that a user can more quickly navigate. For instance, when a given application is going to prompt a user to select between a number of choices, a standard voice prompt might be a full statement of the list of choices followed by a closing voice prompt such as "Please state your response now." In contrast, an expert voice prompt at the same navigation point might be a scaled down version of the list of choices and/or a shorter closing voice prompt such as "Ready." Alternatively, the expert voice prompt might omit portions of the standard voice prompt in their entirety and might substitute a short tone or other mechanism to solicit a response.

10
15
20 As a specific example, at a given navigation point in a weather reporting application, the weather reporting application might define as a standard voice prompt the words, "If you would like to hear today's weather, please say 'today.' If you would like to hear an extended forecast, please say 'extended.' Please speak your response now." And at the same navigation point, the weather reporting application might define an expert voice prompt, "Today or Extended?" followed by a quick prompting tone. Depending on whether the user who is currently interacting with the application is designated as an expert or not, the application may then instruct the platform to apply either the standard prompt or the expert prompt.

In accordance with the exemplary embodiment, a given user can be designated as either a standard-mode user or an expert-mode user. (Other such designations or use-levels are also possible. For instance, a user could be a beginner, intermediate or advanced user.) Further, for a

given user, different use-level designations can be specified for different navigation points. For instance, a user may be designated as a standard-mode user for a particular application (e.g., a particular URI), and the user may be designated as an expert-mode user for another application or for a particular navigation point within an application.

5 Preferably, the user profile store 52 will include an indication, per user, of a use-level, such as whether or not the user is an expert-mode user. If a user is globally designated as an expert mode user for all applications that the voice command platform executes, the user's profile record may include a flag that indicates the user is an expert-mode user. Alternatively, if the user is designated as an expert-mode user for some navigation points and a standard-mode user for others, then the user's profile record can include a table that indicates, per navigation-point, whether the user is an expert-mode user. The table may, for instance, list URIs (and/or other navigation point labels) and, for each URI, provide a flag indicating whether the user is an expert-mode user.

At a given navigation point in a voice command application, the application can then include logic that instructs the processor to apply a standard-mode prompt if the user is a standard-mode user or to apply an expert-mode prompt if the user is an expert-mode user. In response to that logic, the processor may consult the user profile store to determine whether or not the user is an expert-mode user. If the user is an expert-mode user, the processor may then apply the expert-mode prompt. And if the user is a standard-mode user, the processor may apply
20 the standard-mode prompt.

In accordance with the exemplary embodiment, as noted above, the enhanced system logic further includes a mechanism to automatically switch a user from standard mode to expert

mode, either generally or for specific navigation points. The transition to expert mode can occur based on any designated event or combination of events.

For example, as noted above, the processor can be programmed to review a user's navigation history and to automatically transition a user to expert mode after the user has accessed a particular navigation point, particular navigation points, and/or the voice command platform generally, more than a threshold number of times. The processor may perform this review each time a user calls the voice command platform, or at other designated times.

This logic can be as simple or as complex as desired. For instance, the logic could direct the processor to switch a user to expert mode for a given navigation point after the user has accessed the navigation point at least 10 times. Alternatively, the logic could direct the processor to switch a user to expert mode for a given application once the user has accessed the application at least 20 times and has accessed at least two second level navigation points in the application. Other examples are possible as well.

In the exemplary embodiment, when the processor determines that the user should be transitioned to expert mode, the processor could prompt the user with an announcement indicating the benefits of the transition and requesting the user's consent to the transition. Alternatively, the processor could perform the transition without asking the user for consent. Further, the processor could perform the transition even when the user is offline, such as at nighttime when the platform might have more available processing power.

It should be further understood that the processor can also transition a user more granularly through use-modes, rather than just transitioning between two modes (such as a "standard" (not-expert) mode and an "expert" mode). For instance, the expert-mode transition logic might operate to slowly transition a user to more of an expert mode as the user interacts

more with the platform. In this sense, transitioning a user from standard mode to expert mode could be interpreted to mean transitioning the user from any use-mode (which might be considered "standard" to a more advanced use-mode (which might be considered "expert"), even though other modes may exist as well.

5 As an example, the more the processor executes a given set of VXML prompt logic when interacting with a given user, the more abbreviated the prompts may become (as defined by graduated levels of logic in the application, keyed to a use-level set for the user, for instance). The first time, the processor might present a long tutorial prompt to the user. The second time, the processor might present a detailed version of the voice prompt. The 10th time, the processor might present a more abbreviated version of the voice prompt. And the 25th and later times, the processor might present the user with only a tone prompt.

**iii. Using Session Information to Help
Restore a Session After System Disconnect**

As noted above, the enhanced system logic can also facilitate restoring a user to where the user left off in a voice command session, when a session ends abruptly. This can occur, for instance, if the user's telephone connection to the platform gets cut off (e.g., as a result of a coverage drop, signal fade, hang-up or out-dialed call), if the system encounters an error that causes the session to be dropped, or in other instances.

20 In the exemplary embodiment, when a system disconnect occurs, the processor will preferably maintain an indication of the disconnect for a period of time (e.g., approximately 15 minutes) in user profile store 52, together with an indication of the user's navigation points and perhaps other session information at the moment of the disconnect. Provided that the platform

keeps a log of the navigation points that the user has accessed, as described above, the user's latest navigation point could be the most recent navigation point in that log.

When a user calls the platform, the processor may be programmed to consult the user's profile record to determine whether a system disconnect occurred during a session with the user.

5 If the processor determines that a system disconnect occurred, the processor may then prompt the user to indicate whether the user wants to resume where the user left off or, rather, whether the user wants to begin from the start. After the user responds, the processor then either loads and executes an application where the user left off or begins a voice command session normally.

For example, during a session with the voice command platform, a user may be navigating through a set of driving directions. Preferably, the content provider has arranged each step in the driving directions at a separate URI (where one URI leads to the next). When the user navigates to step 5 of the driving directions, the processor may record the URI of step 5 in the user's navigation history, preferably together with a name of the navigation point (e.g., "Step 5 of driving directions"). At that point, a system disconnect might occur. (For instance, the user might intentionally disconnect only to return later to get later steps of the directions.) The processor may then record a flag or other indication in the user profile, indicating that a system disconnect occurred (e.g., in response to the fact that a normal logoff sequence did not occur, or based on any other suitable logic), and noting the time of the disconnect.

20 When the user then calls back the platform, the processor may consult the user's profile and determine that a system disconnect occurred. In response, the processor may programmatically prompt the user to indicate whether the user wants to resume where the user left off. For instance, based on the user's navigation history, the processor may identify the navigation point that the user most recently visited and may then prompt the user accordingly,

such as "Say 'Step 5' to resume Step 5 of the Driving Directions or say 'Voice Command' to start over." Once the user responds, the processor may then either load the associated URI.

d. Providing Unified Message Notification

According to the exemplary embodiment, user profile store 52 can include a consolidated summary of the number of messages waiting for the user at multiple separate message portals, and the processor can be programmed to provide a consolidated or unified message notification to the user. This function can thus involve (i) maintaining consolidated message waiting information for a user and (ii) providing the consolidated notification to the user.

In order to maintain consolidated message waiting information, the platform provider will preferably receive update messages from various message portals (e.g., an e-mail portal, a voice mail portal, a fax portal, an instant messaging portal, etc.), indicating the number of messages waiting for the user at each message portal. To accomplish this, a user may register with various message portals to authorize the message portal to report message-waiting counts to the voice command platform (or to an intermediate entity that functions to store and forward the information to the voice command platform). The various message portals will then send update messages to the voice command platform, autonomously or in response to requests (e.g., periodic requests) from the platform.

Preferably, the update messages that each portal sends to the voice command platform will be in a predetermined standard format that processor 20 is programmed to recognize. For instance, each update message could be an IP-based message that includes an XML indication of parameters such as (i) message portal name, (ii) user ID (e.g., the user's message portal ID or, more preferably, the user's voice command platform ID), and (iii) priority (e.g., standard, urgent, etc.) As an example, a message from a publicly available e-mail portal might specify that the

user has 4 messages of urgent priority waiting at that portal, while a message from a private corporate e-mail portal might specify that the user has 2 messages of standard priority waiting at the corporate e-mail system.

Each portal can be programmed to send such an update message to the voice command platform each time the number of messages waiting for the user at the portal changes and/or upon request from the platform. For instance, once a user reads or hears a message that was waiting at a given message portal, the message portal may decrement the number of messages waiting by one and send an update message to the voice command platform.

In the exemplary embodiment, processor 20 is then programmed to receive the update message count per user from each of multiple message portals and to responsively update the user's profile record with a current summary of the number of messages waiting at the various portals. For instance, the user profile record might include a table that lists various message portals, the type of each portal, and the number of messages (of various priority levels) waiting for the user at the given portal. The processor can thus readily update the table in response to each new update message that it receives from a message portal.

When a user calls the platform, and/or at another designated time during a session with the platform, the processor preferably provides the user with a consolidated message notification, based on the summary maintained in the user profile. The consolidated message notification will indicate to the user how many messages are waiting for the user to retrieve from at least two different message portals. For instance, a system-level root VXML application can welcome the user to the platform and can then announce to the user the number of messages waiting for the user at the various portals. The system-level application may carry out this function before or after providing the user with an initial Ready prompt.

5 The user may then opt to browse to the designated message portals so as to have the messages read to the user. For instance, if the user has messages waiting in a particular e-mail system (e.g., at a given corporate e-mail server), the user may command the platform to load that system's e-mail application (e.g., an application referencing that corporate e-mail server), which may function to read the user's e-mail messages to the user.

10 In the exemplary embodiment, the user profile store 52 can include for each user a respective notification profile record that establishes user-specific features of the consolidated message notification. For example, a user's notification profile might establish user-specific names for each of the user's message portals, and the platform may use those names when presenting the consolidated message notification to the user. For instance, a user's profile record might designate a particular e-mail portal as "work e-mail" and another e-mail portal as "home e-mail." When the platform presents the consolidated notification to that user, the platform may report to the user that the user has "3 messages waiting in your work e-mail account and 8 messages waiting in your home e-mail account" for instance.

15 As another example, a user's notification profile can establish what information the consolidated message notification should include and how the information should be presented. For instance, a user's profile record might indicate that the consolidated notification should summarize messages in categories (such as by level of priority or importance for instance). Further, the user's profile might indicate that the consolidated notification should be filtered in
20 some manner, such as by summarizing only those messages designated as urgent, or only those messages originating from a particular sender, for instance.

Still further, the information provided in the consolidated message notification can vary depending on the type of message portals involved. For instance, if the user has an e-mail

account, a fax account and an instant message account, the consolidated notification could present the user with counts of messages waiting at these three accounts as well as indications of (i) the number of pages included respectively in each fax message and (ii) the identities of the user's instant message "buddies" who are currently online. Other examples are possible as well.

3. Provisioning

In the exemplary embodiment, a provisioning system may be established so as to allow users to set up and manage their profiles. For instance, a provisioning server may be provided on the Internet and may function as an intermediary between the voice command platform and users. As such, the provisioning server may have real-time access to the user profile store 52 and may function to validate data (e.g., user-specified changes to profile records) before writing them to the profile store.

A user may connect with the provisioning server from a web browser on a client computer. Once connected, the user would log in (e.g., by providing a user ID, which the provisioning server would match with a profile record), and the provisioning server would then present the user with a consolidated summary of the user's profile record. The consolidated summary may include information such as: (i) the persona and/or default TTS engine designated for the user, (ii) a list of the user's personal-grammars and their corresponding actual grammars, (iii) the user's expert mode state, (iv) the user's navigation history, (v) a list of the user's bookmarks and (vi) the user's consolidated messaging summary. The provisioning server may then permit the user to modify some or all of this information via the web interface.

In addition, the provisioning server may function as an interface through which content providers can provide the voice command platform with secondary phoneme dictionaries. Preferably, a content provider would log in, and the provisioning server would then allow the

content provider to upload or modify secondary phoneme dictionaries and associations with applications. The provisioning server may then store the secondary dictionaries on the voice command platform or elsewhere to be accessed and used by the platform.

4. Conclusion

5 Exemplary embodiments of the present invention have been described above. Those skilled in the art will understand, however, that changes and modifications may be made to these embodiments without departing from the true scope and spirit of the invention, which is defined by the claims.

09964099-092604